
PCI
SIG®

PCI Express® Advanced Protocol Topics

Debendra Das Sharma, EWG Member

David Harriman, PWG Co-Chair

Copyright © 2005, PCI-SIG, All Rights Reserved

Agenda

- Review: What's New with 1.1
- Updates to PCIe™ Revision 1.1 Base Spec
  - ✓ Errata
  - ✓ New capabilities – ECNs in progress
- Preview of "Gen 2" Development - Draft
  - ✓ Modified speed negotiation mechanism
  - ✓ Compliance mode
- Summary / Call to Action

Copyright © 2005, PCI-SIG, All Rights Reserved

What's New with 1.1

Copyright © 2005, PCI-SIG, All Rights Reserved

New With 1.1

- Hot Plug
  - ✓ Many significant changes for Root/Switch Ports and for Hot Plug System Software to improve hot plug robustness
  - ✓ Implementation simplified for Endpoints
  - ✓ Most changes in Sections 2.2.8.7 & 6.7
- Error Reporting
  - ✓ Improves ability of system software to handle several cases of Completer detected non-fatal errors
  - ✓ Key benefit: UR reporting need not be disabled during device enumeration
  - ✓ See DevCon Presentation “PCIe Error Reporting ECN” for more detail
  - ✓ Most changes in Section 6.2, 7.8.3 & 7.8.4

Copyright © 2005, PCI-SIG, All Rights Reserved

New With 1.1 - Continued

- PME_Turn_Off
  - ✓ All Endpoints must respond to the PME_Turn_Off handshake request
  - ✓ Changes in Sections 5.2 and 5.3
- Flow Control Initialization
  - ✓ Restrictions on Transmitter are loosened
  - ✓ Handshake robustness improved
  - ✓ Changes in Sections 2.6.1 & 3.3
- Reset Limit Adjustment
  - ✓ Time limit from end of Fundamental Reset to entry to LTSSM Detect state changed from 80ms to 20ms
  - ✓ Changes in Section 6.6

Copyright © 2005, PCI-SIG, All Rights Reserved

Rev. 1.1 Base Spec Updates

Copyright © 2005, PCI-SIG, All Rights Reserved

PCIe Rev 1.1 Errata – Class Code Conflict

- Class Code conflict for RC Event Collector
- If you implement RC Event Collector,
  - ✓ Consider flexible mechanisms for setting class code
  - ✓ Wait for errata correcting conflict
- Resolution details still WIP –
  Look for errata to be published

Copyright © 2005, PCI-SIG, All Rights Reserved

Trusted Configuration – New Optional Capability

- *Status: Draft ECN*
- Optional new address space - Trusted Configuration Space (TCS)
  - ✓ Software in Trusted Software Environment may issue Trusted Configuration Requests (new TLPs – see below)
  - ✓ TCS access provided through Trusted Configuration Access Mechanism (TCAM)
- Devices “know” TCS requests “trustworthy”
- Software “knows” device is “trustworthy”
- Provides plumbing mechanism
  - ✓ Additional information required to effectively use
- See DevCon presentation “PCIe Trusted Config Space & Link Speed Controls”

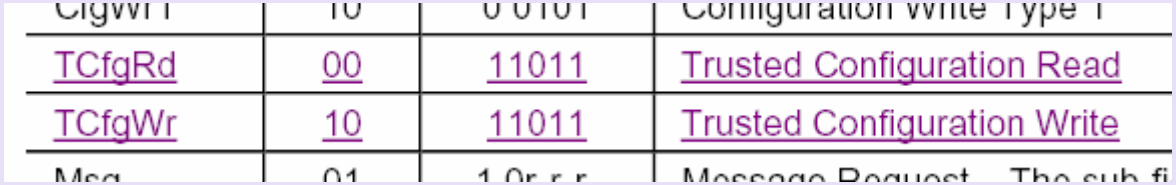

CfgWr1	10	0 0101	Configuration Write Type 1
TCfgRd	00	11011	Trusted Configuration Read
TCfgWr	10	11011	Trusted Configuration Write
Msg	01	1 0r r r	Message Request – The sub-fi

Copyright © 2005, PCI-SIG, All Rights Reserved

PCIe Capability Structure Expansion

- *Status: ECR*
- Adds space for new capabilities
  - ✓ Completion Timeout – see next foil
- Implementation required only as needed
- Adds clear requirements for unimplemented register space

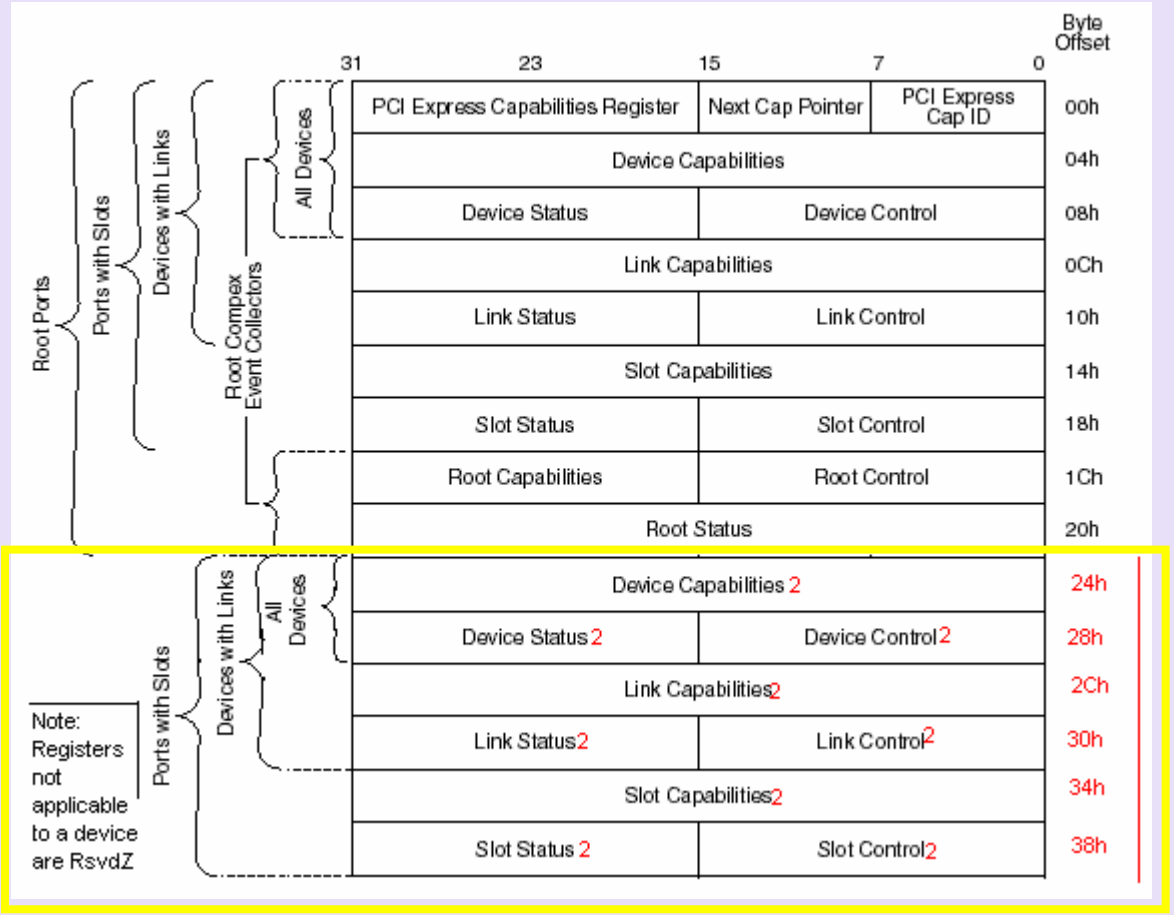

Copyright © 2005, PCI-SIG, All Rights Reserved

Completion Timeout Control

- *Status: ECR*
- Adds capability to disable Completion Timeout
  - ✓ "Required" for all devices implementing ECN
  - ✓ In future, simply required
- Adds *optional* capability for system firmware/software to set Completion Timeout time value
  - ✓ Devices indicate supported ranges from the four bins defined: 50us to 10ms, 10ms to 250ms, 250ms to 4s, 4s to 64s
  - ✓ Two selectable ranges for each bin
    - Example: in 10ms to 250ms range – 16ms to 55ms, 65ms to 210ms
  - ✓ Ranges are separated by invalid value ranges to ensure non-overlapping timeouts

Copyright © 2005, PCI-SIG, All Rights Reserved

Bandwidth Change Notification Mechanism

- *Status: ECR*

- Adds mechanism to signal link has changed width or speed
  - ✓ Uses interrupt & status bit
  - ✓ Covers BW change due to hardware autonomous action or software direction

Copyright © 2005, PCI-SIG, All Rights Reserved

“Gen 2” Updates

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Gen 2: Speed Negotiation

- Link trains to L0 in Gen 1 speed initially.
- Supported speeds advertised in TS sequences
  - ✓ Supported speeds by the other component noted in Config.Complete (and Recovery.RcvrCfg)
- Speed changed through handshake in Recovery state.
  - ✓ Added a new substate Recovery.Speed.
- LTSSM enters Recovery from L0 if a speed change is desired.
  - ✓ Can be initiated by hardware or software.
  - ✓ Both sides exchange speed information (supported speeds as well as intent to change speed) in Recovery.RcvrLock and Recovery.RcvrCfg.
  - ✓ Speed changed in Recovery.Speed.
- Polling.Speed state obsolete.
- Link speed can be changed multiple times.
  - ✓ Reliability and Power improvements.
  - ✓ No need to take the link down
  - ✓ Revert back to Gen 1 if a higher speed is totally inoperable
  - ✓ Revert back to prior speed if a speed change did not succeed.

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Gen 2 LTSSM Changes for Speed Negotiation

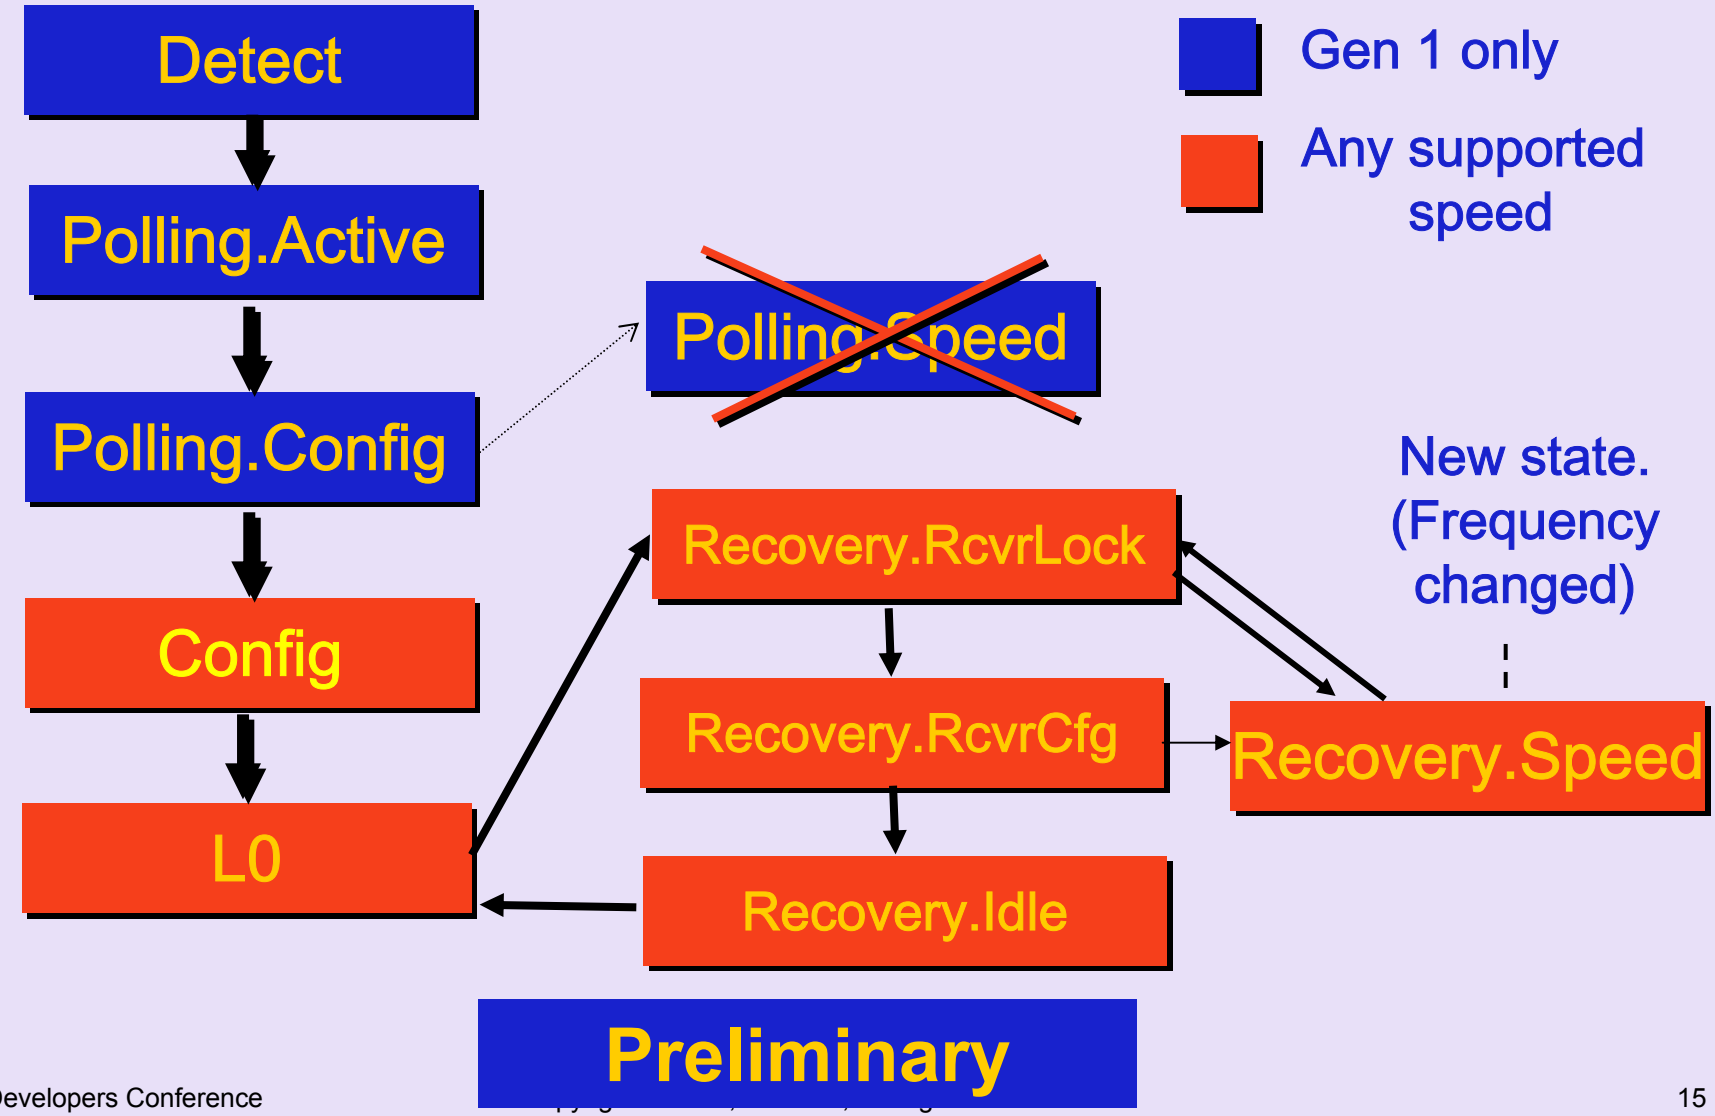

Config Changes for Gen2

- In Link Capabilities:
  Maximum Link Speed → Supported Link Speeds
  - ✓ Reports component capabilities – encoding added for gen2 speed
- New control field: Target Link Speed
  - ✓ Sets target & upper limit on link operational speed
  - ✓ Also sets speed for software initiated Compliance
- New control bit: Hardware Autonomous Speed Control
  - ✓ Controls ability of Hardware to reduce speed

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Config Changes for Gen2 - Continued

- New control bit: Enter Compliance
  - ✓ Software initiated Compliance mode entry
- In Link Status:
  Link Speed → Current Link Speed
  - ✓ Reports the current operational link speed when link is up
- New status bit:
  Link Speed Change Attempt Failed
  - ✓ Set to indicate that hardware was forced to reduce target speed during negotiation

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Speed Negotiation Example – Gen 1 to Gen 2 Change

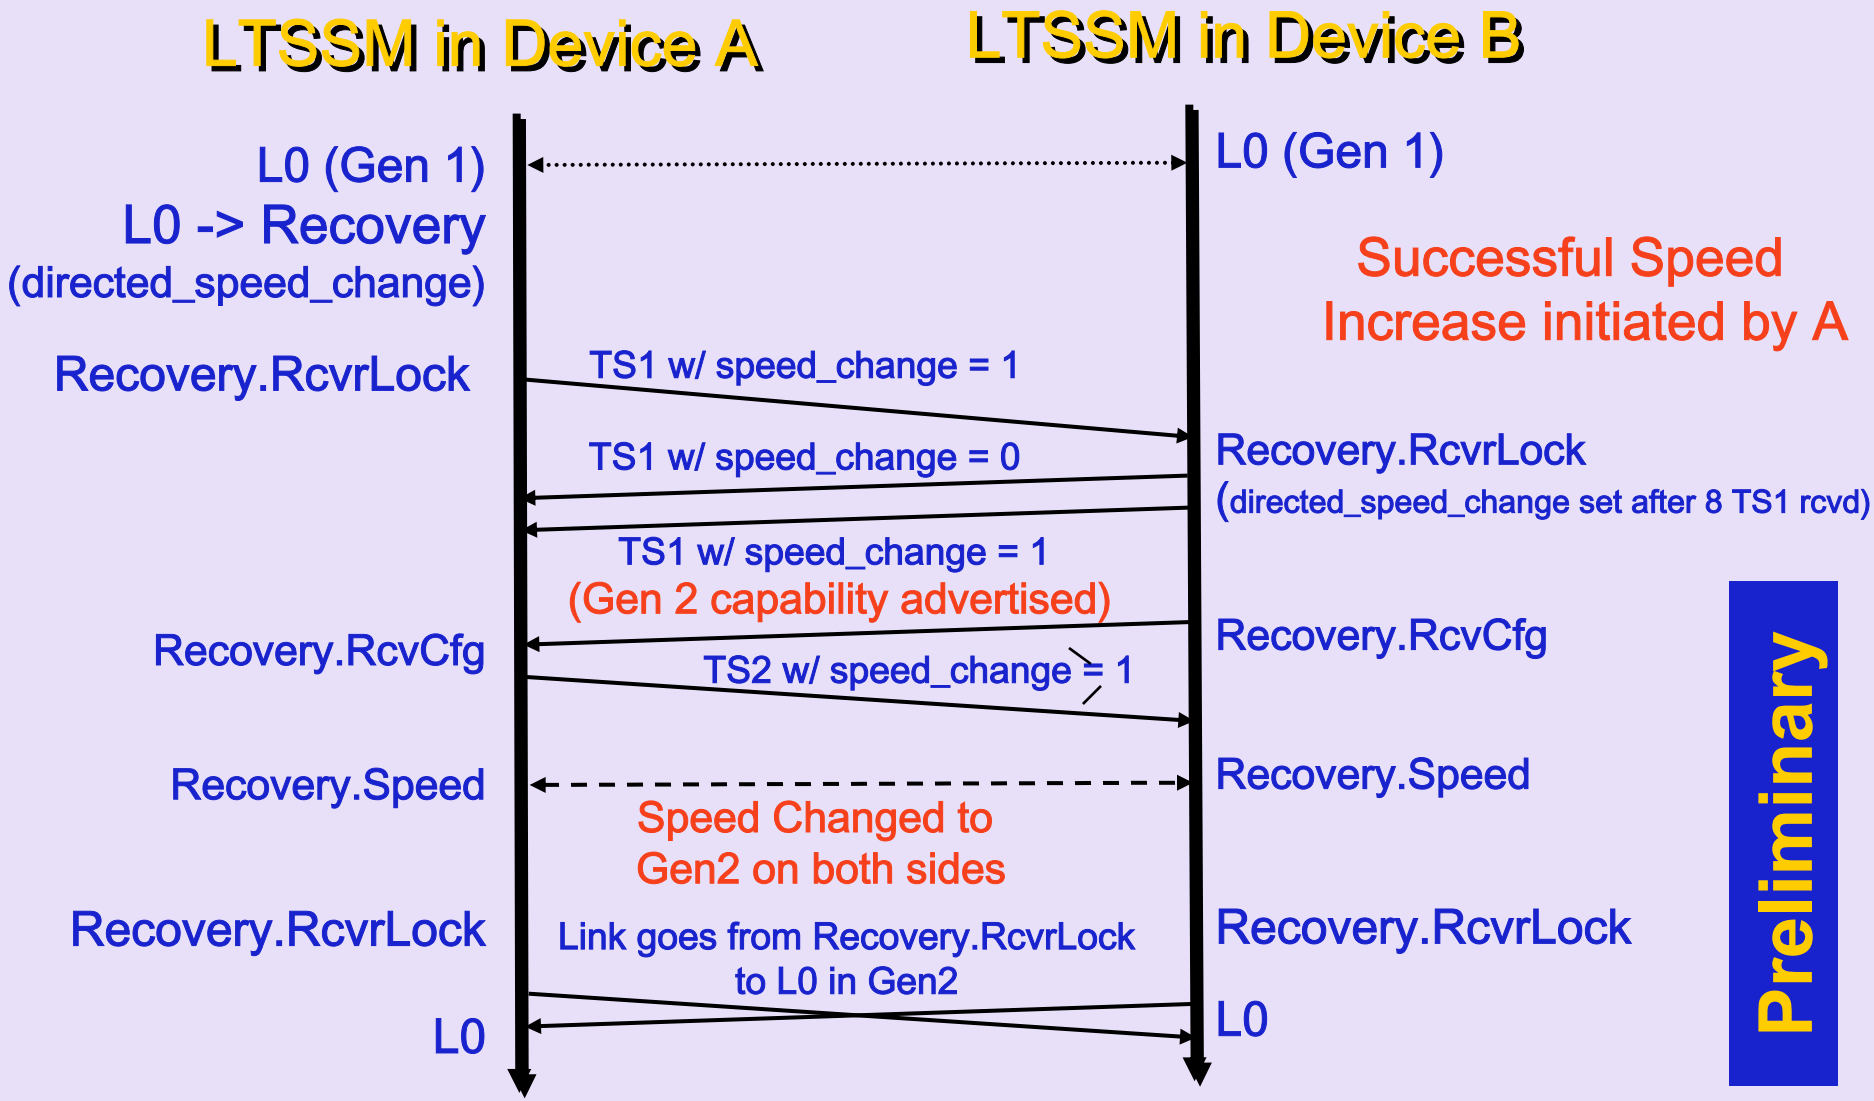

Copyright © 2005, PCI-SIG, All Rights Reserved

Speed Negotiation Example- Gen 2 not functional

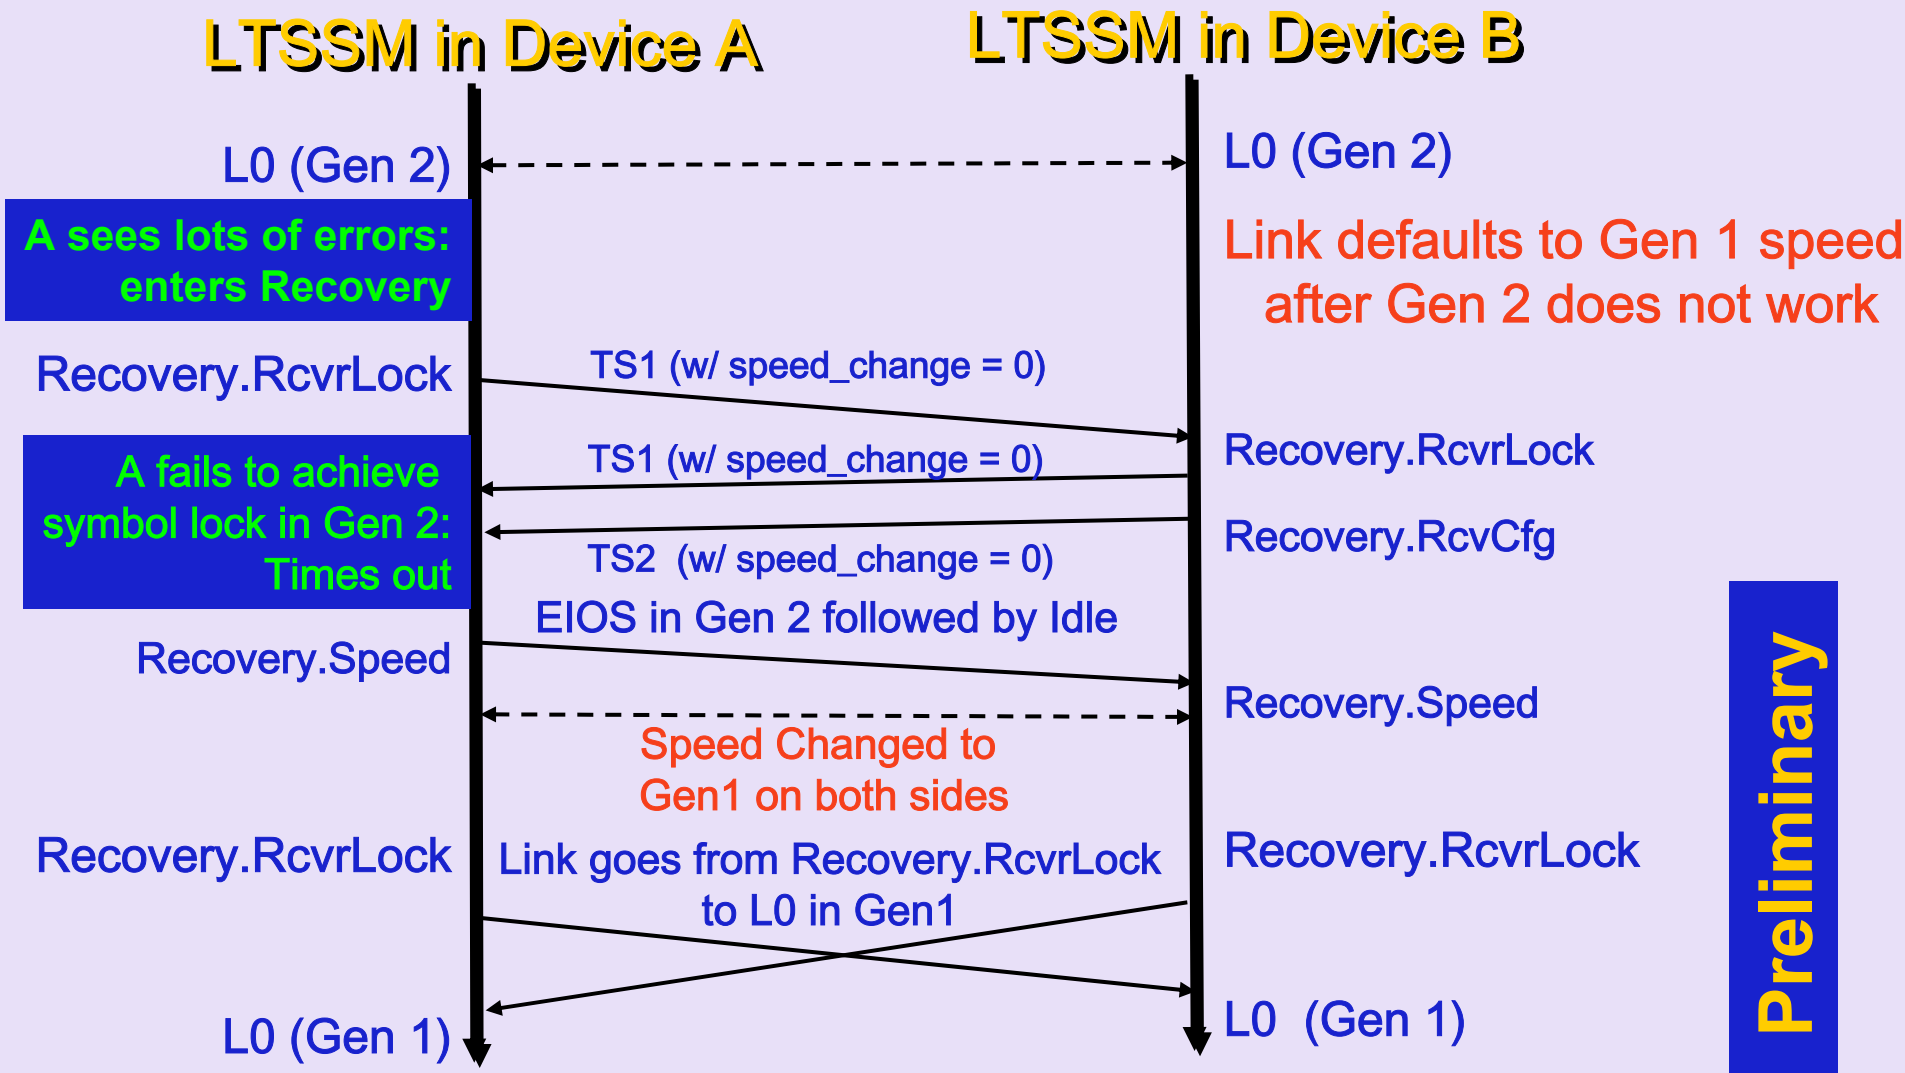

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Gen 2: Polling.Compliance Speed

- Two ways to set the speed
  - ✓ Inband Method:
    - Gen 2 Speed: Tie n and p. Pull up to a DC voltage. Voltage at receiver to be 500-700 mV.
    - Gen 1 Speed: Default. Also can tie n and p and pull down to GND.
    - Need a comparator for determining speed
    - Multilane Link: Speed determined by the highest reported speed among the lanes of the link that has not received an exit from electrical idle.
  - ✓ CSR method (new method)
    - Write to newly defined CSRs to get to compliance at desired speed.
- Speed determined prior to entering Polling.Compliance

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Gen 2: Polling.Compliance Speed

- Polling.Active is always Gen 1
  - ✓ Transition to/from Polling.Compliance operating in Gen 2 speed:
    - – EIOS needs to be transmitted followed by
    - – Electrical idle between 1 ms to 2 ms before changing speed.

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Compliance Speed: Inband

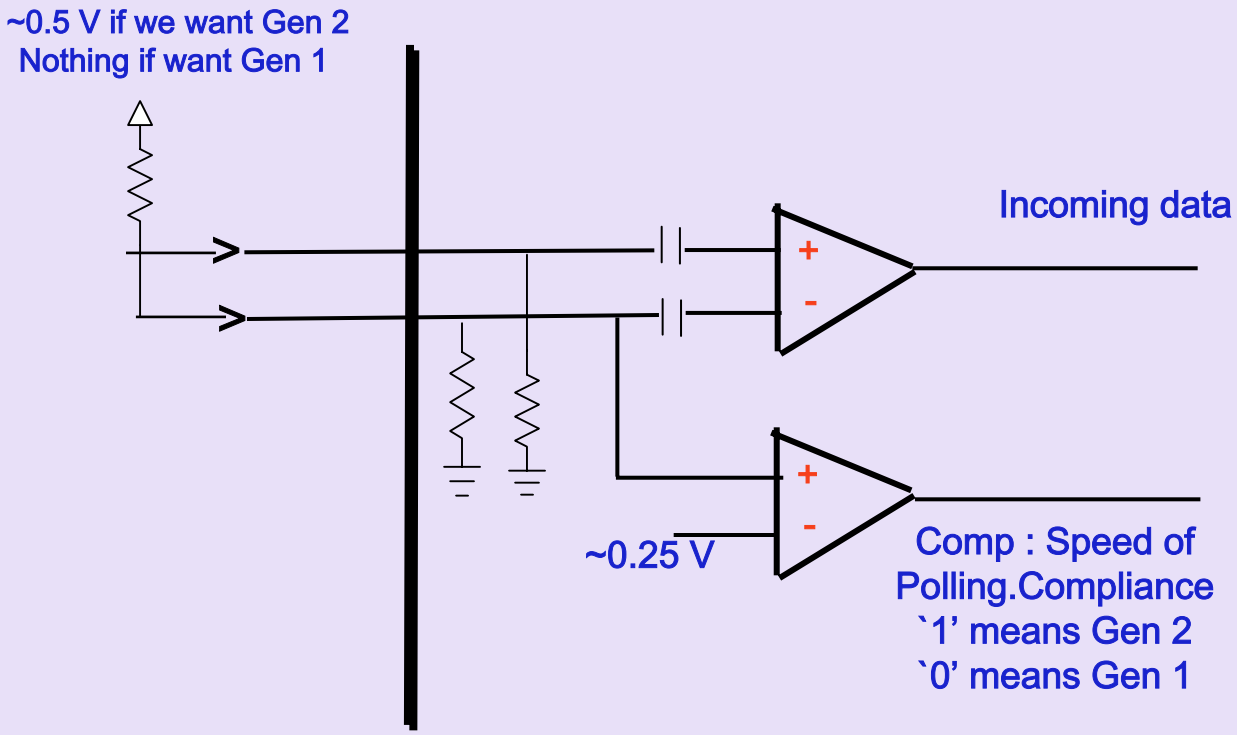

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Other Changes for Gen 2

- Detection of entry to electrical idle a challenge.
  - ✓ Can detect electrical idle using circuits (1.1 method)
  - ✓ Alternatively, can infer entry to electrical idle
    - 128 us of not seeing a COM in L0 (surprise detach)
    - No exit from EI at least once in 128us in Loopback.Active
- High Availability Enhancement w/ Polling.Compliance
  - ✓ Entry condition relaxed from any lane that detected a receiver but did not get an exit from Electrical Idle to multiple lanes.
  - ✓ Must go to Polling.Compliance if all lanes that detected a receiver did not get an exit from electrical idle.
- Electrical Idle Ordered Set extension:
  - ✓ Gen 2 speed: Two consecutive sets of COM, IDL, IDL, IDL sent
  - ✓ Gen 1 speed: one set of COM, IDL, IDL (same as 1.1)
- Electrical Idle detection enhancement (optional):
  - ✓ Received signals switching at a frequency greater than 125MHz

Preliminary

Copyright © 2005, PCI-SIG, All Rights Reserved

Summary & Call to Action

- Base all new components on Rev. 1.1
  - ✓ Many improvements over the 1.0a
- Comprehend errata and additions
  - ✓ Errata comprehension required for compliance
  - ✓ ECN implementation permitted but not required
- Track Gen 2 development work to be ready for future speed increase
  - ✓ 1.1 Components not directly affected

Copyright © 2005, PCI-SIG, All Rights Reserved

Thank you for attending the
PCI-SIG Developers Conference 2005.

For more information please go to
www.pcisig.com

Copyright © 2005, PCI-SIG, All Rights Reserved

PCI Express Advanced Protocol Topics

Debendra Das Sharma, EWG Member

David Harriman, PWG Co-Chair

Copyright © 2005, PCI-SIG, All Rights Reserved


PCI
SIG®